arXiv:0901.0876v1 [stat.CO] 7 Jan 2009

A Fast Algorithm for Robust Regression with Penalised Trimmed Squares

L. Pitsoulis and G. Zioutas
Department of Mathematical and Physical Sciences
Aristotle University of Thessaloniki
541 24 Thessaloniki, Greece

January 7, 2009

Abstract

The presence of groups containing high leverage outliers makes linear regression a difficult problem due to the masking effect. The available high breakdown estimators based on Least Trimmed Squares often do not succeed in detecting masked high leverage outliers in finite samples.

An alternative to the LTS estimator, called Penalised Trimmed Squares (PTS) estimator, was introduced by the authors in [28, 29] and it appears to be less sensitive to the masking problem. This estimator is defined by a Quadratic Mixed Integer Programming (QMIP) problem, where in the objective function a penalty cost for each observation is included which serves as an upper bound on the residual error for any feasible regression line. Since the PTS does not require presetting the number of outliers to delete from the data set, it has better efficiency with respect to other estimators. However, due to the high computational complexity of the resulting QMIP problem, exact solutions for moderately large regression problems is infeasible.

In this paper we further establish the theoretical properties of the PTS estimator, such as high breakdown and efficiency, and propose an approximate algorithm called Fast-PTS to compute the PTS estimator for large data sets efficiently. Extensive computational experiments on sets of benchmark instances with varying degrees of outlier contamination, indicate that the proposed algorithm performs well in identifying groups of high leverage outliers in reasonable computational time.

Keywords: Robust regression; quadratic mixed integer programming; least trimmed squares; outliers detection.

Contents

1 Introduction **2**

2 Trimmed Squares Regression **4**
2.1 Preliminaries 4
2.2 Least Trimmed Squares 5

3 Penalised Trimmed Squares **6**
3.1 Properties of the PTS 8
3.1.1 Exact Fit Property and Breakdown Point 9
3.1.2 Efficiency 11
3.2 Penalties Computation 11
3.2.1 Robust scale estimate $\hat{\sigma}$ 11
3.2.2 Unmasking Multiple High Leverage Outliers 12
3.3 Reinclusion of Observations 13
3.4 Quadratic Programming Formulation 13

4 Fast-PTS **14**
4.1 The Algorithm 16

5 Computational Results **18**
5.1 Exact Solutions 19
5.2 Benchmark Instances 20
5.3 Monte Carlo Simulation 21

6 Conclusions and Future Research **24**

1 Introduction

In multilinear regression models experimental data often contains outliers and bad influential observations, due to errors. It is important to identify these observations and eliminate them from the data set, since they can lead the regression estimate to take erroneous values. If the data is contaminated with a single or few outliers the problem of identifying such observations is not difficult, but in most cases data sets contain groups of outliers which makes the problem more difficult due to masking and swamping effects. An indirect approach to outlier identification is through a robust regression estimate. If a robust estimate is relatively unaffected from outliers, then the residuals from the robust fit should be used to identify the outliers.

Two commonly used criteria for robustness are the *breakdown point* and *efficiency* of the estimator. The breakdown point can be roughly defined as the minimum percentage of outliers present in the data, that could lead the error between the robust estimate and the hypothetical true estimate to be infinitely large. It is desirable for a robust estimator to have a breakdown point of close to 50% of the sample size.

A well known estimator with high breakdown point is the Least Trimmed Squares (LTS) estimator of Rousseeuw and Leroy [22]. The method consists of finding a $n - k$ subset of observations whose deletion from the data set would lead to the k smallest residual sum of squares. However, it is well known that

the LTS loses efficiency. Several robust estimators have been proposed in the literature as extensions to the LTS, to obtain high breakdown point and simultaneously improve the efficiency. Among them are the S-estimator proposed by Rousseeuw and Yohai [23], the MM-estimator proposed by Yohai [26], the $r-$estimator proposed by Yohai and Zamar [27], the S1S estimator proposed by Coakley and Hettmansperger [6], WLSE computed from an initial high breakdown robust estimator by Agostinelli and Markatou [1] and the REWLS by Gervini and Yohai [9]. Most of these estimators start from an initial robust scale σ_{LTS} and robust regression coefficient estimate $\boldsymbol{\beta}_{LTS}$ given from the LTS estimator with coverage $k \simeq [(n+p+1)/2] \simeq 50\%$. Then they compute the standardised residuals using different schemes based on robust scale σ_{LTS} and design weights with appropriate cut-off values, and they reconsider which of the potential outliers should remain in the basic subset as clean data and which should be eliminated as outliers. Thus, they simultaneously attain the maximum breakdown point of the LTS estimator while improving its efficiency. The key to the success of all the aforementioned methods is to start with a good initial regression coefficient estimate $\boldsymbol{\beta}_{LTS}$. When a finite sample contains multiple high leverage points, these may be masked outliers and would affect the LTS estimator resulting in a biased initial estimate. Moreover, the LTS method requires a priory knowledge of the *coverage* k, or equivalently the number $n-k$ of the most likely outliers that produces the largest reduction in the residual sum of squares when deleted. Unfortunately, this knowledge of k is typically unknown, Gentleman and Wilk [8]. Computation of the LTS estimator requires the solution of a hard combinatorial problem, and there have been many exact and approximation algorithms proposed in the literature (see Section 2.2).

A different approach to obtain a robust estimate has been proposed in [29] called Penalised Trimmed Squares (PTS), which does not require presetting the number $n-k$ of outliers to delete from the data set. The PTS approach "trims" outliers from the data but instead of discarding a fixed number of observations, a fixed threshold for the allowable size of the adjusted residuals is used. The new estimator PTS is defined by minimising a *loss function*, which is the sum of squared residuals and penalty costs for deleting bad observations. These penalties regulate the threshold of the allowable adjusted residuals, as well as the coverage. In order to overcome the problem of groups of masking outliers containing almost identical high leverage points, lower penalties are proposed yielding a smaller adjusted residual threshold for such observations. These penalties are a function of robust leverages resulted from the MCD estimator by Rousseeuw and Van Driessen [20]. Computationally, the PTS estimator also involves the solution of a hard combinatorial problem.

The purpose of this paper is twofold. First, the robust properties of the PTS estimator as presented in [29] such as equivariance, exact fit property, high breakdown point and efficiency, are established. Secondly we present an efficient approximation algorithm called Fast-PTS, to compute the PTS estimator for large data sets without having to solve the problem exactly.

The organisation of the paper is the following. Section 2 contains some preliminary notations and definitions. The definition of the PTS estimator, its

robust properties, the computation of the penalties and finally an equivalent mixed integer quadratic programming formulation are presented in Section 3. Section 4 describes the Fast-PTS algorithm, where initially a set of necessary optimality conditions is given, and the algorithm is then presented in pseudocode. Extensive computational experiments to examine the performance of the Fast-PTS algorithm with respect to, both the quality of the approximate solution when compared to the exact solution of small instances, and the robust performance to a set of benchmark and artificially generated large instances, are given in Section 5. Finally, concluding remarks and further possible extensions to both the algorithm and the robust estimator are given in the last section.

2 Trimmed Squares Regression

2.1 Preliminaries

In this section we will state some well established facts in order to set up the notation that will be used for the rest of the paper. We consider the multi-linear regression model with p independent variables

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{u}, \tag{1}$$

where $\mathbf{y} = (y_1, y_2, \ldots, y_n)^T$ is the $n \times 1$ vector of the response variable, $\mathbf{X}$ is a full rank $n \times p$ matrix with rows $\mathbf{x}_i = (x_{i1}, x_{i2}, \ldots, x_{ip})$ of explanatory variables, $\boldsymbol{\beta}$ is a $p \times 1$ vector $\boldsymbol{\beta} = (\beta_1, \beta_2, \ldots, \beta_p)^T$ of unknown parameters , and $\mathbf{u}$ is a $n \times 1$ vector $\mathbf{u} = (u_1, u_2, \ldots, u_n)^T$ of iid random errors with expectation zero and variance σ^2. We observe a sample $(x_{i1}, x_{i2}, \ldots, x_{ip}, y_i)$, for $i = 1, 2, \ldots, n$, and construct an estimator for the unknown parameter $\boldsymbol{\beta}$. The Ordinary Least Squares Estimator (OLS) is defined by minimising the squared residual loss function

$$OLS(\mathbf{X}, \mathbf{y}) := \arg\min_{\boldsymbol{\beta}} \sum_{i=1}^{n} r(\boldsymbol{\beta})_i^2 \tag{2}$$

where $r(\boldsymbol{\beta})_i$ is the regression residual $r(\boldsymbol{\beta})_i := y_i - \mathbf{x}_i\boldsymbol{\beta}$. We will write r_i instead of $r(\boldsymbol{\beta})_i$ whenever the parameter vector $\boldsymbol{\beta}$ need not be explicitely stated. It is well known that a solution to (2) is obtained in polynomial time (i.e. $\mathcal{O}(n^3)$) by the normal equations

$$OLS(\mathbf{X}, \mathbf{y}) = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{y}$$

A transformation of the residuals that will be useful in this work is the adjusted residual α_i which is defined as

$$\alpha_i := \frac{r_i}{\sqrt{1 - h_i}}, \tag{3}$$

where $h_i \quad (0 < h_i < 1)$ measures the leverage of the i^{th} observation and is defined as $h_i := \mathbf{x}_i^T(\mathbf{X}^T\mathbf{X})^{-1}\mathbf{x}_i$. The reduction in the sum of squared residuals

in (2) due to the deletion of an observation $(\mathbf{x}_i, y_i)$ is the square of its adjusted residual

$$\alpha_i^2 = \frac{r_i^2}{1 - h_i}.$$

Unfortunately, points that are far from the predicted line (outliers) are overemphasised. There are several types of outliers that can occur in a data set $(\mathbf{X}, \mathbf{y})$. Following the terminology of Rousseeuw and Van Zomeren [24], a point $(\mathbf{x}_i, y_i)$ which does not follow the linear pattern of the majority of the data but whose $\mathbf{x}_i$ is not outlying is called a *vertical outlier*. A point $(\mathbf{x}_i, y_i)$ whose $\mathbf{x}_i$ is outlying (large h_i) is called a *leverage point*. It is called *good* leverage point when $(\mathbf{x}_i, y_i)$ follows the pattern of the majority otherwise it is called *bad* leverage point. The OLS estimator is very sensitive to outliers, in the sense that the presence of any type of previously mentioned types of outliers greatly affects the solution of (2). We wish to construct a robust estimator for the parameter $\boldsymbol{\beta}$, such that the influence of any observation $(\mathbf{x}_i, y_i)$ on the sample estimator is bounded.

2.2 Least Trimmed Squares

Rousseeuw introduced the Least Trimmed Squares (LTS) estimator in [19], which fits the best subset of k observations, removing the rest $n - k$ observations. The LTS estimator is defined as follows defined as:

$$\begin{aligned} LTS(\mathbf{X}, \mathbf{y}, k) := \arg\min_{\boldsymbol{\beta}} \quad & \sum_{i=1}^{k} r(\boldsymbol{\beta})_{(i)}^2 \\ \text{s.t.} \quad & r(\boldsymbol{\beta})_{(1)}^2 \leq r(\boldsymbol{\beta})_{(2)}^2 \leq \cdots \leq r(\boldsymbol{\beta})_{(n)}^2 \end{aligned} \tag{4}$$

where k is called the *coverage*, and is chosen as $k \geq [(n + p + 1)/2]$ a priori, so as to maximise the breakdown point. In (4) and in what follows, we employ the convention of writing $r(\boldsymbol{\beta})_{(i)}$ for the i^{th} smallest residual error with respect to $\boldsymbol{\beta}$. The LTS estimator has high breakdown point but loses efficiency, since $n-k$ observations have to be removed from the sample even if they are not outliers. Assuming that the set of points $(\mathbf{x}_i, y_i)$ are in *general position*, that is any two distinct subsets of p points have different regression lines, and observing that for a given k the problem in (4) consists of solving $\binom{n}{k}$ ordinary least squares problems, it is immediate that it has a unique solution.

Problem (4) has been formulated as a nonlinear programming problem with linear constraints and non-convex objective function by Giloni and Padberg in [10], thereby its local minima are well characterised by the Karush-Kuhn-Tucker optimality conditions. In particular the authors in [10] establish its equivalency with a concave minimisation problem, which is known to be NP-Hard, and provide a procedure for computing local minima. The exact computation of (4) requires an exponential number of steps with respect to the size of the problem, therefore it can only be applied to small instances. Agulló in [2] presented a branch and bound procedure for the exact computation of (4) along with several procedures to reduce the computational effort, and managed to find

the global optimum for several benchmark instances of size $n < 50$. Polynomial time algorithms for the exact computation of the LTS for simple ($p = 1$) linear regression problems are given by Hössjer [15], and recently by Li [16] who utilises advanced data structures to attain improved computational complexity. Due to the high computational complexity of the LTS estimator there have been a plethora of approximation algorithms appeared in the literature with varying success. We mention among others, the Forward Search algorithm by Atkinson [3, 4], the Feasible Solution algorithm by Hawkins [12, 14] and the more recent Fast-LTS algorithm by Rousseeuw and Van Driessen [21] which currently the most accurate and efficient algorithm.

3 Penalised Trimmed Squares

The basic idea of the PTS estimator is to insert fixed penalty costs for each observation into the objective function of (4), such that only observations whose adjusted residuals is larger than their penalty costs are deleted from the data set. Therefore instead of defining a priori a coverage k, the penalty costs are defined and the coverage becomes a decision variable. In the sections that follow suitable penalties for multiple high-leverage outliers are proposed. We consider as most likely outliers that subset of the observations that produces significant reduction in the sum of squared residuals when deleted. The proposed PTS estimator can be defined through a minimisation problem, with an objective function split into two parts; the sum of k squared residuals in the clean data and the sum of the penalties for deleting the rest $n - k$ observations.

$$\begin{aligned} PTS(\mathbf{X}, \mathbf{y}, \mathbf{p}) := \arg\min_{\boldsymbol{\beta}, k} \quad & \sum_{i=1}^{k} r(\boldsymbol{\beta})^2_{(i)} + \sum_{i=k+1}^{n} p_{(i)} \\ \text{s.t.} \quad & r(\boldsymbol{\beta})^2_{(1)} \leq r(\boldsymbol{\beta})^2_{(2)} \leq \cdots \leq r(\boldsymbol{\beta})^2_{(n)} \end{aligned} \tag{5}$$

where $\mathbf{p} = (p_1, \ldots, p_n)$ and p_i is the *penalty* cost for deleting the i^{th} observation defined as

$$p_i := \max\{\epsilon, (c\hat{\sigma})^2\}$$

for some small $\epsilon > 0$, where $\hat{\sigma}$ is a *robust residual scale* and c is the *cut-off parameter*. Although these penalties are constant in the basic definition of the PTS, as we shall see later in section 3.2 each observation will have an individual penalty. The estimator performance is very sensitive to the penalties defined a priori which regulate the robustness and the efficiency of the estimator. The choice of the robust scale $\hat{\sigma}$ plays an important role in the coverage of the PTS estimator. The minimisation problem (5) is not convex since it is equivalent to a quadratic mixed integer programming problem, nevertheless there exists a unique solution under mild assumptions on the data (see section 3.4).

Let us define with $f_{PTS}(\boldsymbol{\beta}, k)$ the objective function of (5). We have the following straightforward observations regarding the optimisation problem in (5), where the set of feasible solutions is all $(\boldsymbol{\beta}, k)$ for $\boldsymbol{\beta} \in \mathbb{R}^p$ and $k = 0, \ldots, n$:

i) For any feasible solution $(\boldsymbol{\beta}, k)$ there exists an associated set of k observations as determined by the ordering

$$r(\boldsymbol{\beta})^2_{(1)} \leq r(\boldsymbol{\beta})^2_{(2)} \leq \cdots \leq r(\boldsymbol{\beta})^2_{(k)} \leq \cdots \leq r(\boldsymbol{\beta})^2_{(n)}.$$

ii) For any $\boldsymbol{\beta} \in \mathbb{R}^p$ there exists k^* such that

$$f_{PTS}(\boldsymbol{\beta}, k^*) \leq f_{PTS}(\boldsymbol{\beta}, k), \ \ \forall k = 0, \ldots, n,$$

where k^* is determined by those observations with squared residuals w.r.t. $\boldsymbol{\beta}$ less than their respective penalties. Therefore at optimality, $\boldsymbol{\beta} \in \mathbb{R}^p$ determines k uniquely.

iii) Since the second term of the objective function in (5) is independent of $\boldsymbol{\beta}$, for any feasible solution $(\boldsymbol{\beta}, k)$

$$f_{PTS}(\boldsymbol{\beta}_{OLS}, k) \leq f_{PTS}(\boldsymbol{\beta}, k),$$

where $\boldsymbol{\beta}_{OLS}$ is the ordinary least squares estimate obtained for the set of observations so defined by $(\boldsymbol{\beta}, k)$.

We summarise the above in the following lemma which will be used later.

Lemma 3.1. *If $(\boldsymbol{\beta}_{PTS}, k_{PTS})$ is the solution obtained by (5) for sample data $(\mathbf{X}, \mathbf{y})$ then*

$$OLS(\mathbf{X}_{k_{PTS}}, \mathbf{y}_{k_{PTS}}) = \boldsymbol{\beta}_{PTS},$$

where $(\mathbf{X}_{k_{PTS}}, \mathbf{y}_{k_{PTS}})$ is the submatrix and subvector of $(\mathbf{X}, \mathbf{y})$ respectively, as determined by the set of k_{PTS} observations defined by $(\boldsymbol{\beta}_{PTS}, k_{PTS})$.

In view of the above lemma and (3) we can now state the **general principle** of the PTS estimator, that is to delete an observation i if its squared adjusted residual is larger its the penalty cost

$$\alpha_i^2 > (c\hat{\sigma})^2. \tag{6}$$

Therefore, the adjusted residual has as a threshold the square root of the deleting penalty $c\hat{\sigma}$. Given that the PTS estimate is the OLS of the clean data set, the resulted adjusted residuals must be smaller or equal than the penalties

$$\frac{|r_i|}{\sqrt{1 - h_i^*}} < c\hat{\sigma}, \ \ 0 \leq i \leq k. \tag{7}$$

where h_i^* is the leverage of the i^{th} observation in the k observations. Otherwise, the i^{th} observation can not remain in the coverage sample due to the PTS principal.

The PTS as defined in (5) "trims" outliers from the data but instead of discarding a fixed number of observations, a fixed threshold $(c\hat{\sigma})$ for the allowable size of the adjusted residuals is used.

3.1 Properties of the PTS

Most of the robust properties of the PTS estimator are inherited by the use of the robust scale $\hat{\sigma}$, which in our case will be obtained from the LTS estimator with coverage $(n+p+1)/2$. In what follows we will employ the notation used in [22].

Let $T(\mathbf{X}, \mathbf{y})$ be an estimator for some sample data $(\mathbf{X}, \mathbf{y})$. An estimator T is called *regression equivariant* if

$$T(\mathbf{X}, \mathbf{y} + \mathbf{X}\mathbf{v}) = T(\mathbf{X}, \mathbf{y}) + \mathbf{v}, \quad \text{for any } \mathbf{v} \in \mathbb{R}^n,$$

scale equivariant if

$$T(\mathbf{X}, c\mathbf{y}) = cT(\mathbf{X}, \mathbf{y}), \quad \text{for any } c \in \mathbb{R},$$

and *affine equivariant* if

$$T(\mathbf{X}\mathbf{A}, \mathbf{y}) = \mathbf{A}^{-1}T(\mathbf{X}, \mathbf{y}), \quad \text{for any nonsingular } \mathbf{A} \in \mathbb{R}^{p\times p}.$$

Lemma 3.2. *The PTS estimator is regression, scale and affine equivariant.*

Proof. For regression equivariance, by Lemma 3.1 there will exist some subset of k observations were

$$\begin{aligned} PTS(\mathbf{X}, \mathbf{y} + \mathbf{X}\mathbf{v}, \mathbf{p}) &= (\mathbf{X}_k^T\mathbf{X}_k)^{-1}\mathbf{X}_k^T(\mathbf{y}_k + \mathbf{X}_k\mathbf{v}) \\ &= (\mathbf{X}_k^T\mathbf{X}_k)^{-1}\mathbf{X}_k^T\mathbf{y}_k + \mathbf{v} \\ &= PTS(\mathbf{X}, \mathbf{y}, \mathbf{p}) + \mathbf{v}. \end{aligned}$$

Similarly for scale and affine equivariance. □

The PTS estimator determines that subset of observations whose deletion produces the largest reduction in the sum of squared residuals. The LTS estimator determines the subset of size k with the minimum sum of squared residuals. The relationship between the two estimators can be seen by the following propositions.

Proposition 3.1. *If the PTS estimator for sample data* $(\mathbf{X}, \mathbf{y})$ *converges to the solution* $(\boldsymbol{\beta}_{PTS}, k_{PTS})$*, then* $LTS(\mathbf{X}, \mathbf{y}, k_{PTS}) = \boldsymbol{\beta}_{PTS}$*.*

Proof. Observe that any k observations will have the same sum of penalties, therefore we will have

$$\begin{aligned} LTS(\mathbf{X}, \mathbf{y}, k_{PTS}) &= \arg\min_{\boldsymbol{\beta}} \sum_{i=1}^{k_{PTS}} r(\boldsymbol{\beta})^2_{(i)} \\ &= \arg\min_{\boldsymbol{\beta}} \left(\sum_{i=1}^{k_{PTS}} r(\boldsymbol{\beta})^2_{(i)} + (n - k_{PTS})c\hat{\sigma} \right) \\ &= \boldsymbol{\beta}_{PTS} \end{aligned}$$

□

Conversely under mild assumptions, from the LTS estimator given in (4) and using the largest residual r_k as the deleting penalty, the PTS estimator given in (5) leads to the same results.

Proposition 3.2. *If for sample data* $(\mathbf{X}, \mathbf{y})$ *and coverage* k *we have* $LTS(\mathbf{X}, \mathbf{y}, k) = \boldsymbol{\beta}_{LTS}$ *while the largest residual* $r(\boldsymbol{\beta}_{LTS})_k$ *is an increasing function with respect to* k*, then for fixed penalty* $p_i = (r(\boldsymbol{\beta}_{LTS})_k)^2, \quad i = 1, \ldots, n$*, we have* $PTS(\mathbf{X}, \mathbf{y}, \mathbf{p}) = \boldsymbol{\beta}_{LTS}$.

Proof. Let $PTS(\mathbf{X}, \mathbf{y}, \mathbf{p}) = (\boldsymbol{\beta}_{PTS}, k_{PTS})$. By Proposition 3.1 enough to show that $k_{PTS} = k$. Moreover the largest residual $r(\boldsymbol{\beta}_{PTS})_{k_{PTS}}$ will also be an increasing function with respect to k_{PTS}. We will have for any $k \geq k_{PTS}$

$$\begin{aligned} \sum_{i=1}^{k_{PTS}} r_i^2 + \sum_{i=k_{PTS}}^{n} p_i &= \sum_{i=1}^{k} r_i^2 - \sum_{i=k_{PTS}}^{k} r_i^2 + \sum_{i=k}^{n} p_i + \sum_{i=k_{PTS}}^{k} p_i \\ &= \sum_{i=1}^{k} r_i^2 + \sum_{i=k}^{n} p_i + \sum_{i=k_{PTS}}^{k} (p_i - r_i^2) \end{aligned}$$

where the terms in the last summation are all nonnegative since $p_i \geq (r(\boldsymbol{\beta}_{LTS})_i)^2$, $i = k_{PTS}, \ldots, k$. Therefore the minimum is obtained for $k = k_{PTS}$. Similarly for $k \leq k_{PTS}$. □

In the LTS estimator the coverage k has to be chosen between $n/2$ and n. For coverage $k = [(n + p + 1)/2]$ the maximum breakdown point is obtained by loosing efficiency whereas for $k = n$ the obtained breakdown value is reduced to $1/n$. As it is the case with most robust estimators, there is a trade-off between robustness and efficiency. The advantage of the PTS procedure is that there is no need to define a priory the coverage k. For given penalty $(c\hat{\sigma})^2$, where $\hat{\sigma}$ is a high breakdown robust scale, the PTS can resist against all influential observations with large adjusted residuals (i.e. $\alpha_i > c\hat{\sigma}$).

3.1.1 Exact Fit Property and Breakdown Point

In this section we will investigate two common robust properties for the PTS estimator, namely the exact fit and breakdown points. The *exact fit point* of an estimator T is defined as the minimum fraction of observations on a hyperplane, necessary to guarantee that the estimator coincides with that hyperplane [27]. Formally

$$\delta^*(T, \mathbf{X}, \mathbf{y}) \quad := \quad \min \left\{ \frac{m}{n} : \exists \boldsymbol{\beta} \text{ such that } \mathbf{X}_m \boldsymbol{\beta} = \mathbf{y}_m, T(\mathbf{X}, \mathbf{y}) = \boldsymbol{\beta} \right\}, \quad (8)$$

where $(\mathbf{X}_m, \mathbf{y}_m)$ is the submatrix and subvector of $(\mathbf{X}, \mathbf{y})$ respectively, as determined by the set of m observations. A dual notion is used in [22], where the exact fit point is taken as the minimum fraction of observations *not* on a hyperplane that will force the estimator to give a different estimate from that hyperplane. An estimator is said to have the *exact fit property* when $\delta^* \leq 0.5$.

The *breakdown point* of an estimator T is defined as the minimum fraction of observations that could take arbitrary values and force the estimator to vary indefinitely from its original estimate. This is the finite version of breakdown point as it was introduced in [7], and could be stated formally as

$$\epsilon^*(T, \mathbf{X}, \mathbf{y}) \quad := \quad \min\left\{ \frac{m}{n} : \sup_{\mathbf{X}_m, \mathbf{y}_m} \|T(\mathbf{X}, \mathbf{y}) - T(\mathbf{X}_m, \mathbf{y}_m)\| = \infty \right\}, \quad (9)$$

where $(\mathbf{X}_m, \mathbf{y}_m)$ is the submatrix and subvector of $(\mathbf{X}, \mathbf{y})$ respectively, as determined by the set of m observations. The relationship between exact fit point and breakdown point for regression equivariant estimators is

$$\delta^* \geq 1 - \epsilon^* \quad (10)$$

where under mild assumptions equality also holds.

Proposition 3.3. *If for sample data $(\mathbf{X}, \mathbf{y})$ there exists some $\boldsymbol{\beta}$ such that strictly more than $\frac{n+p-1}{2}$ of the observations satisfy $\mathbf{x}_i\boldsymbol{\beta} = y_i$ and are in general position, then the PTS solution with robust scale σ_{LTS} equals $\boldsymbol{\beta}$.*

Proof. Since the above is true for the LTS estimator (see corollary in page 134 of [22]), we will have $\sigma_{LTS} = 0$. Therefore the penalties as defined in (5) will be $\mathbf{p} = (\epsilon, \ldots, \epsilon)$ and at optimality it will hold

$$PTS(\mathbf{X}, \mathbf{y}, \mathbf{p}) = LTS(\mathbf{X}, \mathbf{y}, k_{LTS}) = \boldsymbol{\beta},$$

for $k_{LTS} = \frac{n+p+1}{2}$. □

It is known [22] that for regression and scale equivariant estimators it holds

$$\epsilon^* \leq \frac{n - p + 2}{2n}.$$

Below we show that the PTS estimator has a high breakdown point.

Proposition 3.4. *The breakdown point of the PTS estimator with robust scale σ_{LTS} is bounded from below as*

$$\epsilon^*_{PTS} \geq \frac{n - p - 1}{2n}.$$

Proof. By Lemma 3.2 and (10) we know that

$$\delta^*_{PTS} \geq 1 - \epsilon^*_{PTS},$$

while by Proposition 3.3 we get

$$\delta^*_{PTS} \leq \frac{n + p + 1}{2n}.$$

Therefore by the above

$$\epsilon^*_{PTS} \geq 1 - \frac{n + p + 1}{2n} = \frac{n + p - 1}{2n}.$$

□

3.1.2 Efficiency

Let $(\mathbf{X}, \mathbf{y})$ be sample data where the random errors $\mathbf{u}$ as given in the model (1) follow the normal distribution, $\mathbf{u} \sim N(0, \sigma^2)$.

Consider now that $n \to \infty$. We know from [22] that the variance σ_{LTS} obtained from the LTS estimator is consistent (i.e. $\lim_{n\to\infty} \sigma_{LTS} = \sigma$). Moreover for each observation i we have that the leverage $h_i = 0$, which means that by the PTS general principle an observation i will be deleted if

$$a_i^2 = r(\boldsymbol{\beta}_{OLS})_i^2 > c\hat{\sigma}^2$$

Therefore, under Gaussian conditions on the random errors and for large n, we can say that the PTS estimator with robust scale $\hat{\sigma} = \sigma_{LTS}$ and cut-off parameter $c \in [2, 3]$, deletes a very small fraction of observations as outliers.

3.2 Penalties Computation

3.2.1 Robust scale estimate $\hat{\sigma}$

The key to the success of the PTS estimator is to use a robust scale $\hat{\sigma}$ for evaluating the proper penalties. This robust scale $\hat{\sigma}$ is obtained via the LTS estimator with coverage $k_{LTS} = [(n+p+1)/2]$ as follows. Initially a preliminary error scale $\hat{s}$ is computed

$$\hat{s} = c_{k_{LTS},n} \sqrt{\frac{1}{k_{LTS}} \sum_{i=1}^{k_{LTS}} r(\boldsymbol{\beta}_{LTS})_{(i)}^2},$$

where the constant $c_{k_{LTS},n}$ is so chosen in order to make the scale estimation consistent at the Gaussian model, that is

$$c_{k_{LTS},n} = 1/\sqrt{1 - \frac{2n}{k_{LTS}\alpha_{k_{LTS},n}} \Phi\left(\frac{1}{\alpha_{k_{LTS},n}}\right)},$$

where

$$\alpha_{k_{LTS},n} = \frac{1}{\Phi^{-1}(\frac{k_{LTS}+n}{2n})}.$$

Then we improve the efficiency results of the LTS by applying a re weighted procedure (see [22]). First the standarized residuals $r(\boldsymbol{\beta}_{LTS})_i/\hat{s}$ are computed, and they are used to determine a weight for each observation as follows:

$$w_i = \begin{cases} 1 & \text{if } \frac{r(\boldsymbol{\beta}_{LTS})_i}{\hat{s}} |\hat{s}| \leq 2.5, \\ 0 & \text{otherwise.} \end{cases}$$

The final robust scale is then given by

$$\hat{\sigma} = \sqrt{\frac{\sum_{i=1}^n w_i r(\boldsymbol{\beta}_{LTS})_i^2}{\sum_{i=1}^n w_i - p}}. \tag{11}$$

3.2.2 Unmasking Multiple High Leverage Outliers

Based on the PTS principal, an observation is deleted if its adjusted residual is greater than the square root of the penalty cost $\alpha_i \geq c\hat{\sigma}$. For y-outliers and even for some $\mathbf{x}$-outliers such an approach has successful performance. Unfortunately the masking problem arises when there is a group of high leverage points in the same direction. As it is known, the leverage value h_i can be distorted by the presence of a collection of points, which individually have small leverages but collectively form a high leverage group. Peña and Yohai [17] point out that the individual leverage h_i of each point might be small $h_i << 1$, whereas the final residual may appear to be very close to 0, and this is a masking problem. Thus in a set of identical high leverage outliers, the adjusted residual is masked and might be too small $\alpha_i << c_i\hat{\sigma}$, and as a consequence a masked high leverage outlier may not be deleted.

In order to overcome the distortion caused by the masking problem and unmask the leverage h_i of the potential high leverage points, we are based on the Minimum Covariance Determinant (MCD) procedure (Rousseeuw and Van Driessen [20]). For given coverage k, the MCD procedure yields the "clean" $\mathbf{X}_k$ matrix where $n-k$ vectors $\mathbf{x}_i$ corresponding to high leverage observations have been removed from the original matrix $\mathbf{X}$. Thus, for coverage k close to 51% (i.e. $k = [(n+p+1)/2]$), the leverage of each point $(\mathbf{x}_i, y_i)$, $i = k+1, \ldots, n$ as it joins the clean data subset taken from MCD is

$$h_i^* = \mathbf{x}_i^T (\mathbf{X}_{k,i}^T \mathbf{X}_{k,i})^{-1} \mathbf{x}_i, \quad \text{for } \; i = k+1, \ldots, n,$$

where $\mathbf{X}_{k,i}$ is the matrix $\mathbf{X}_k$ plus the row $\mathbf{x}_i$. In the above h_i^* is the unmasked leverage of each one of the potential high leverage points which can be considered as the leverage at the clean data set yielded from the MCD procedure. For the rest of the observations $(\mathbf{x}_i, y_i)$ the new leverages now are

$$h_i^* = \mathbf{x}_i^T (\mathbf{X}_k^T \mathbf{X}_k)^{-1} \mathbf{x}_i, \quad \text{for } \; i = 1, \ldots, k.$$

Given that the high leverage identical outliers have been removed from the clean matrix $\mathbf{X}_k$ due to the MCD procedure, the new leverage values h_i^* of the masked points will appear larger than the original h_i. Thus, in order to overcome the masking problem we transform the adjusted residuals as follows:

$$\alpha_i^* = \frac{r_i}{\sqrt{1-h_i}} \cdot \frac{\sqrt{1-h_i}}{\sqrt{1-h_i^*}} = \frac{r_i}{\sqrt{1-h_i^*}}. \tag{12}$$

Based on the reliable results of the MCD procedure, the new robust adjusted residual in (12) is a good diagnostic criterion for all observations even for the multiple high leverage outliers. Incorporating the new robust adjusted residual now in the PTS estimator, we can see how the penalties get individualised for each observation since an observation now will be considered as an outlier if

$$\alpha_i^* > (c\hat{\sigma}) \Rightarrow r_i > c\sqrt{1-h_i^*}\hat{\sigma},$$

that is the penalty for each observation is defined as

$$p_i := (c\sqrt{1-h_i^*}\hat{\sigma})^2 \tag{13}$$

where c is the cut-off parameter, $\hat{\sigma}$ the robust scale while h_i^* the robust leverage.

Thus, the PTS estimator tends to remove the masked high leverage observations from the data set, unless their residuals are very small. However, according to the above analysis, good high leverage points may also appear as extreme points with large leverages h_i^* and the yielded adjusted residual threshold may be too small. As a consequence, some of the good high leverage points may be deleted from the data set. This affects only the efficiency of the solution. However, the efficiency of the final estimate is improved by testing each potential outlier one by one in the reinclusion stage of the PTS procedure which is described in the next section.

3.3 Reinclusion of Observations

In order to improve the efficiency of the $\hat{\sigma}_k$ all observations with small studentized residuals are reincluded into the clean subset of size k. We follow the reinclusion test for each of the $n-k$ observations, similarly to Hadi and Simonoff [11] and Pena and Yohai [17]. All the $n-k$ points deleted initially stage are tested one by one using the studentized predicted error for outlyingness

$$t_i = \frac{y_i - \mathbf{x}_i^T \boldsymbol{\beta}_{PTS}}{\hat{\sigma}\sqrt{1+h_i}}, \quad k+1 \leq i \leq n,$$

where $h_i = \mathbf{x}_i^T[\mathbf{X}_k^T\mathbf{X}_k]^{-1}\mathbf{x}_i$. Under normality assumptions, t_i follows Student's t distribution with k degrees of freedom. Therefore, every observation $(\mathbf{x}_i, y_i)$, is reincluded in the subset of the clean data if $t_i \leq t_{\alpha/2,k}$. We have found empirically that for small samples, $t_{\alpha/2,k} = 2$ works well. The final PTS estimate is the OLS solution of the resulting data set after the reinclusion.

3.4 Quadratic Programming Formulation

The PTS estimator as defined in (5) can be stated as a Quadratic Mixed Integer Programming (QMIP) problem, by expanding the residuals and adding 0-1 decision variables, as follows:

$$\begin{aligned}
\min_{\boldsymbol{\beta}} \quad & \sum_{i=1}^{n} \left(r_i^2 + \delta_i (c\hat{\sigma})^2\right) \\
\text{s.t.} \quad & \mathbf{x}_i^T\boldsymbol{\beta} + r_i \geq y_i - s_i \\
& \mathbf{x}_i^T\boldsymbol{\beta} - r_i \leq y_i + s_i \\
& s_i \leq \delta_i M \\
& \delta_i \in \{0,1\} \\
& r_i, s_i \geq 0 \quad i = 1, \ldots, n,
\end{aligned} \tag{14}$$

where s_i can be regarded as the pulling distance for moving an outlier towards the regression line, δ_i is a decision variable which takes the value of 1 if observation i is to be removed from the clean data and 0 otherwise, and M is an upper bound on the residuals $r_i, i = 1, \ldots, n$. Given any fixed $\boldsymbol{\delta} \in \{0,1\}^n$, from the 2^n possible ones, and using matrix notation problem (14) becomes

$$\begin{aligned} \min_{\boldsymbol{\beta}} \quad & \mathbf{r}^T\mathbf{r} + \boldsymbol{\delta}^T\mathbf{p} \\ \text{s.t.} \quad & \mathbf{X}\boldsymbol{\beta} + \mathbf{r} \geq \mathbf{y} - \mathbf{s} \\ & \mathbf{X}\boldsymbol{\beta} - \mathbf{r} \leq \mathbf{y} + \mathbf{s} \\ & \mathbf{s} \leq \boldsymbol{\delta} M \\ & \mathbf{r}, \mathbf{s} \geq \mathbf{0}, \end{aligned}$$

where $\mathbf{p} = ((c\hat{\sigma})^2, (c\hat{\sigma})^2, \ldots, (c\hat{\sigma})^2)^T$, $\mathbf{r} = (r_1, \ldots, r_n)^T$, $\mathbf{s} = (s_1, \ldots, s_n)^T$, $\mathbf{y} = (y_1, \ldots, y_n)^T$, $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_n)^T$ and the matrix $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n]^T$. This problem has linear constraints and a convex quadratic objective function, since the Hessian of $\mathbf{r}^T\mathbf{r}$ has nonnegative eigenvalues (and it is therefore positive semi-definite). Therefore we have a convex program, which will have a unique global optimum solution according to the Karush-Kuhn-Tucker optimality conditions [5]. Considering that there is a finite number of possible $\boldsymbol{\delta}$, we can conclude that under the assumption that the data are in general position, problem (14) or equivalently (5) has a unique solution.

Due to the high computational complexity of the resulting QMIP problem the PTS estimator as defined by (14) cannot be solved exactly in reasonable time, even for moderately sized problem instances (i.e. $n < 50$). This is because the computational time required by even the state of the art integer solvers such as CPLEX, grows exponentially with respect to the size of the problem instance. Moreover it is almost certain (unless $P = NP$) that there does not exist an exact polynomial time algorithm that solves (14). Therefore the need arises for a specialised approximation algorithm, which we present in the next section.

4 Fast-PTS

Let us provide a combinatorial formulation of problem (14) in order to facilitate the discussion of the algorithm. Observe that $\delta_i = 1$ implies $r_i = 0$, while for $\delta_i = 0$ we have that $r_i = r(\boldsymbol{\beta})_i$ which is the residual error of observation i with respect to the regression vector $\boldsymbol{\beta}$. If we let the set of observations be O we could equivalently formulate problem (14) as follows:

$$\min_{T \subseteq O} \qquad L(T) := \sum_{i \in T} r(\boldsymbol{\beta}_T)_i^2 + \sum_{i \in O/T} p_i \tag{15}$$

where p_i is the associated penalty for observation i defined in (13), and $\boldsymbol{\beta}_T$ is the OLS solution for the set T of observations uniquely defined as

$$\boldsymbol{\beta}_T := (\mathbf{X}_T^T\mathbf{X}_T)^{-1}\mathbf{X}_T^T\mathbf{y}_T, \tag{16}$$

where $\mathbf{X}_T, \mathbf{y}_T$ are submatrices of $\mathbf{X}$ and $\mathbf{y}$ respectively, whose rows are indexed by the observations in T. The equivalence of problem (15) with (14) stems from the one-to-one correspondence between $\boldsymbol{\beta}_T$ and $T \subseteq O$, as it is defined by the normal equations in (16). The following necessary conditions for optimality can be stated.

Proposition 4.1. *If $T \subseteq O$ is an optimal solution to (15) then*

$$\begin{aligned} &i) \quad r(\boldsymbol{\beta}_T)_i^2 \leq p_i, \ \ \forall i \in T \\ &ii) \quad r(\boldsymbol{\beta}_T)_i^2 \geq p_i, \ \ \forall i \in O/T \end{aligned}$$

Proof. For the case *i)* let us define the set

$$S_1 := \{i \in T : r(\boldsymbol{\beta}_T)_i^2 > p_i\}$$

and assume that $S_1 \neq \emptyset$. Then

$$\begin{aligned} L(T) = \sum_{i \in T} r(\boldsymbol{\beta}_T)_i^2 + \sum_{i \in O/T} p_i \quad &> \quad \sum_{i \in T/S_1} r(\boldsymbol{\beta}_T)_i^2 + \sum_{i \in O/\{T/S_1\}} p_i \\ &\geq \quad \sum_{i \in T/S_1} r(\boldsymbol{\beta}_{T/S_1})_i^2 + \sum_{i \in O/\{T/S_1\}} p_i \end{aligned}$$

therefore we have found a set T/S_1 were $L(T) > L(T/S_1)$ which is a contradiction to the hypothesis. For the case *ii)* analogously define the set

$$S_2 := \{i \in O/T : r(\boldsymbol{\beta}_T)_i^2 < p_i\}$$

and assume that $S_2 \neq \emptyset$. Then

$$\begin{aligned} L(T) = \sum_{i \in T} r(\boldsymbol{\beta}_T)_i^2 + \sum_{i \in O/T} p_i \quad &> \quad \sum_{i \in T} r(\boldsymbol{\beta}_T)_i^2 + \sum_{i \in S_2} r(\boldsymbol{\beta}_T)_i^2 \sum_{i \in O/\{T \cup S_2\}} p_i \\ &\geq \quad \sum_{i \in T \cup S_2} r(\boldsymbol{\beta}_{T \cup S_2})_i^2 + \sum_{i \in O/\{T \cup S_2\}} p_i \end{aligned}$$

which implies that $L(T) > L(T \cup S_2)$ and again it contradicts the original hypothesis of the optimality of T. Therefore $S_1 = S_2 = \emptyset$ and the proposition is proved. □

It can be easily verified by a counterexample that the above conditions are not sufficient. Nevertheless the number of local optima which they characterise seems to be much smaller than that characterised by more classic k-exchange neighbourhoods which grow exponentially with respect to the size of the associated QMIP. Specifically, if we define the *k-exchange neighbourhood* of a feasible solution T to (15) as

$$N(T)_k := \{S \subseteq O : |T \Delta S| \leq k\}$$

```
procedure Fast-PTS(X, y, p, n, p, MaxIter, RandomSeed)
1    T_min := {1, ..., n};
2    do k = 1, ..., MaxIter →
3         T := Construction(X, y, p, n, p, RandomSeed);
4         T := LocalSearch(X, y, p, n, p, T);
5         if (L(T) < L(T_min)) → T_min := T
6    od;
7    return(β_{T_min})
end Fast-PTS;
```

Figure 1: The Fast-PTS algorithm

where Δ stands for the symmetric difference of sets, then we have the trivial necessary conditions for optimality

$$T \text{ optimal to (15)} \;\Rightarrow\; L(T) \leq L(S),\ \forall S \in N(T)_k. \tag{17}$$

A local search algorithm based on (17) will have to have to search the k neighbourhood of a feasible solution entirely, which will require an exponential number of steps on k with either a breadth-first or a depth-first search strategy. Therefore even for small $k = 1, 2$ this type of local search is computationally very expensive.

The first necessary optimality condition of Proposition 4.1 is used to construct a feasible solution in a randomised fashion, while the second is used to obtain the local optimum in the region defined by this solution.

4.1 The Algorithm

The Fast-PTS algorithm is presented in pseudocode in Figure 1. It accepts as inputs the problem instance data, and the parameters `MaxIter` and `RandomSeed` which specify the maximum number of iterations and a seed for the random number generator respectively. Initially the best solution is T_{min} is set to be all the observations, while this solution is updated in line 5 of Figure 1. Each iteration of the Fast-PTS algorithm is composed of two phases, a construction phase where a good solution is built in a greedy randomised fashion starting from an empty solution, and an improvement phase where the solution from the previous phase is improved to ensure local optimality. At the end of each iteration an approximate solution to (15) is provided. The maximum number of iterations is provided by the user, and the best solution found among all iterations is returned by the algorithm as the approximation to the optimum solution. This randomised procedure for constructing a solution before performing local search, has substantial experimental evidence of good performance for NP-Hard combinatorial optimisation problems (see [18]) since it helps the algorithm to examine a wider area of the feasible space without getting entrapped in local optima.

The construction phase of the algorithm is shown in Figure 2. Let us call a partial solution of (15) say $T \subseteq O$ *penalty free* if

$$r(\boldsymbol{\beta}_T)_i^2 < p_i, \ \ \forall i \in T.$$

In the construction phase of the algorithm, initially a partial solution $T \subset O$ is constructed by choosing at random $(p+1)$ observations such that T is penalty free (line 1 in Figure 2). Note that any p observations are penalty free since they have zero residuals i.e. we have an exact fit. Then this solution is enlarged in a greedy randomised fashion maintaining its penalty free property. Given any partial solution T the following set of possible candidates is defined

$$C(T) := \{j \in O/T : r(\boldsymbol{\beta}_{T\cup j})_i^2 < p_i, \ \ \forall i \in T \cup j\}$$

The elements of $C(T)$ are then sorted in ascending order according to their objective function values as given in (7), that is

$$L(T \cup j_1) \leq L(T \cup j_2) \leq \cdots \leq L(T \cup j_{|C(T)|}),$$

and one observation s is chosen randomly among the first $\max\{1, \alpha|C(T)|\}$ candidates, and the partial solution is updated (i.e. $T := T \cup \{s\}$). The parameter $\alpha \in [0, 1]$ controls the degree of greediness versus randomness in the construction of the solution. If $\alpha = 1$ the the algorithm constructs the solution in a randomised fashion while if $\alpha = 0$ the solution is purely greedy. This procedure is repeated until we reach a partial solution T with $C(T) = \emptyset$, where the construction phase ends.

In lines 2-16 of the `Construction` procedure shown in Figure 2 the main iterations upon which a solution T_c is constructed, are shown. The iterations in lines 4-10 compute the set $C(T_c)$ of candidate elements to be added in the solution, where the penalty free property is maintained in line 5. The cost as well as the index of the associated candidate element are added into a heap in line 8, which will be later used for sorting these values. Finally in lines 11-15, an element j_c is chosen at random among the best candidates, and added into our current solution. The procedure `outheap()` returns the index j of the observation with the smallest $L(T_c \cup j)$ value as computed in and inserted into the heap in lines 7 and 8.

The solution T_c computed by the construction phase is penalty free, and we also know that there is no $j \in O/T_c$ such that $T_c \cup \{j\}$ is penalty free with respect to $\boldsymbol{\beta}_{T_c\cup\{j\}}$. However there may be an observation $j \in O/T_c$ such that its residual with respect to $\boldsymbol{\beta}_{T_c}$ is less than its respective penalty. That is, the necessary optimality conditions stated in Proposition 4.1 may not be satisfied. Therefore further improvement of the solution provided by the construction phase could be achieved. This is performed in the improvement phase of the Fast-PTS algorithm, where the solution by the construction phase is changed to satisfy both necessary optimality conditions. Given a solution T_1 and its associated regressor vector $\boldsymbol{\beta}_{T_1}$, a new solution T_2 is computed as

$$T_2 := \{i \in O : r(\boldsymbol{\beta}_{T_1})_i^2 < p_i\},$$

```
procedure Construction(X, y, p, n, p, RandomSeed)
1    T_c := ConstructRandom(X, y, p, n, p, RandomSeed);
2    while (C(T_c) ≠ ∅) →
3        C(T_c) := ∅;
4        do j ∈ O/T_c →
5            if (r(β_{T_c∪j})_i^2 < p_i,  ∀i ∈ T_c ∪ j) →
6                C(T_c) := C(T_c) ∪ j;
7                L(T_c ∪ j) := Σ_{i∈T_c∪j} r(β_{T_c∪j})_i^2 + Σ_{i∈O/T_c∪j} p_i;
8                inheap(L(T_c ∪ j), j);
9            endif;
10       od;
11       t =random[1, α|C(T_c)|];
12       do j = 1, ..., t →
13           j_c := outheap();
14       od;
15       T_c := T_c ∪ j_c ;
16   endwhile;
17   return(T_c)
end Construction;
```

Figure 2: The `Construction` procedure of the Fast-PTS algorithm

and this is repeated until $T_1 = T_2$. As it is shown in the proof of Proposition 4.1, at each iteration $L(T_1) \geq L(T_2)$ therefore the procedure will terminate in a finite number of steps (usually the convergence is in the order of 4 to 5 steps). This simple local search procedure is depicted in Figure 3.

Overall the computational complexity of the Fast-PTS algorithm is polynomial on the size of the problem, since for a fixed number of iterations, each iteration requires a polynomial number of solutions to an OLS problem, which is itself polynomially solvable.

5 Computational Results

In this section we present the results of the computational experiments performed in order to verify the performance of the Fast-PTS algorithm, and the overall robust behaviour of the proposed estimator. First we compare the solutions found by the Fast-PTS algorithm for the QMIP problem (14) with the exact solutions found by ILOG's CPLEX software, on a set of artificial problem instances. Then we compare the robust behaviour of the algorithm with two of the most succesfull algorithms in robust regression i.e. the Fast-LTS algorithm as implemented by Rousseeuw and Van Driessen [21] and the Fast-S algorithm as presented by Barrera and Yohai [25]. These comparisons were performed on the following sets of instances:

```
procedure LocalSearch(X, y, p, n, p, T)
1     T_1 := T;
2     while(T_2 ≠ T_1) →
3           T_1 := T_2;
4           T_2 := ∅;
5           do i = 1, ..., n →
6                 if (r(β_{T_1})_i^2 < p_i) →
7                       T_2 := T_2 ∪ i;
8                 endif;
9           od;
10    endwhile;
11    return(T_1);
end LocalSearch;
```

Figure 3: The `LocalSearch` procedure of the Fast-PTS algorithm

- a set of "benchmark" data sets to regression outlier problems widely known in the literature,
- a set of Monte Carlo generated data sets, proposed by Barrera and Yohai in [25],

The Fast-PTS algorithm was implemented on Intel's Fortran Compiler version 10.1 for Linux with the optimisation flags `-ipo, -O3, -no-prec-div, -static`, and ILOG's CPLEX integer programming solver version 9.1 was used. All computational experiments where performed on a Intel Core Duo 2.4GHz computer with 2GB of memory running openSUSE Linux 10.3[1].

5.1 Exact Solutions

In Table 1 we can see the performance of the Fast-PTS algorithm with respect to exact solutions for 8 problem instances. In each instance the problem was solved exactly by the CPLEX mixed integer quadratic programming solver, and the Fast-PTS obtained the exact solution for all instances. The number of maximum iterations was set to 100, although in all cases the optimal solution was obtained in fewer iterations. As we can see from Table 1, the computational effort required for the exact solution grows exponentially with respect to the size of the instance, while the Fast-PTS algorithm requires a fraction of a second. Similar comparison would be desirable for larger size instances such as the ones presented in section 5.3. This however is infeasible since the required CPU time would be excessively large for the exact solution, and the size of the branch and cut tree generated by CPLEX exceeds the available memory capacity. Specifically for the problem set in Table 1, for instances with $n \geq 128$, CPLEX required

[1]The code is available upon request at `pitsouli@gen.auth.gr`

larger memory for the tree and terminated the solution process.

data set	n	p	Exact solution	CPLEX (CPU seconds)	Fast-PTS (CPU seconds)
PTS1	48	2	12642.9	6.05	0.038
PTS2	58	2	25910.7	5.68	0.048
PTS3	68	2	17658.1	146.4	0.064
PTS4	78	2	26158.8	42.3	0.076
PTS5	88	2	37951.5	1368	0.096
PTS6	98	2	41289.2	2256	0.092
PTS7	108	2	48519.3	32055	0.132
PTS8	118	2	52829.8	37948	0.152

Table 1: Comparison of Fast-PTS with CPLEX

5.2 Benchmark Instances

Five widely used benchmark instances were examined in this experimental study. The first four are taken from Rousseeuw and Leroy [22] and have been studied by many statisticians in robust literature. The last one is an artificial example from Hadi and Simonoff [11]. Table 2 gives the corresponding results for the PTS estimator, indicating the name of the data set, its dimension and number of observations, the included outliers, the percentage of outliers that have been identified and the running times in seconds. The penalties for deleting outliers are $(2\sqrt{1-h_i^*}\cdot\hat{\sigma})^2$, as have been proposed in section 2. These results are similar to those reported in the literature for other high breakdown estimators and they indicate that the proposed method PTS behaves reliably on these test sets.

Telephone Data. The first data set contains the total number of telephone calls made in Belgium during the years 1950-1973. During the period 1964-1969, (cases 15-20), another recording system was used, hence these observations are unusually high while cases 14 and 21 are marginal. The outliers draw the OLS regression line upwards, masking the true outliers, while swamping in the clean cases 22-24 as too low. The high breakdown estimators Fast-LTS and Fast-S correctly flags the outliers. Also, our PTS estimator correctly identify the true outliers.

Hertzsprung-Russell Stars Data. This data set contains 47 stars in the direction of Cygnus, where x is the logarithm of the effective temperature at the surface of a star and y is the logarithm of its light intensity. Four stars, called giants, have low temperature with high light intensity. These outliers are cases 11, 20, 30 and 34 which can be clearly seen by a scatter plot. The three high breakdown methods identify successfully the outliers.

Modified Wood Gravity Data. This is a real data set with five independent variables. It consists of 20 cases; some of them were replaced by Rousseeuw [19] to contaminate the data by few outliers, namely cases 4, 6, 8 and 19. Once again, all three methods manage to reveal the outliers.

Hawkins, Bradu and Kass Data. The data have been generated by Hawkins et. al [13] for illustrating the merits of a robust technique. This artificial data set offers the advantage that at least the position of the good or bad leverage points is known. The Hawkins, Bradu and Kass data consists of 75 observations in four dimensions. The first ten observations form a group of identical bad leverage points, the next four points are good leverage while the remaining are good data. The problem in this case is to fit a hyperplane to the observed data. Plotting the regression residuals from the model obtained from the standard OLS estimator, the bad high-leverage point data are masked and do not show up from the residual plot. Some robust methods not only fail to identify the outliers, but they also swamp in the good cases 11-14. The Fast-LTS and Fast-S estimators correctly flag the outliers. For the PTS estimator, the MCD procedure reveals the first 14 points of this data set as high leverage points. Application of the PTS to these data, starting with robust scale estimate about $\hat{\sigma} = 0.61$ and downweighting the penalty cost with penalties from (13), rejects only the first 10 points as outliers, which are known as the bad leverage points.

Hadi Data. These data have been created by Hadi and Simonoff [11], in order to illustrate the performance of various robust methods in outlier identification. The two predictors were originally created as uniform $(0, 15)$ and were then transformed to have a correlation of 0.5. The depended variable was then created to be consistent with the model $y = x_1 + x_2 + u$ with $u \sim N(0, 1)$. The first 3 cases $(1 - 3)$ were contaminated to have predictor values around $(15, 15)$, and to satisfy $y = x_1 + x_2 + 4$. Scatterplots or diagnostics have failed to detect the outliers. Many identification methods fail to identify the three outliers. Some bounded influence estimates have largest absolute residual at the clean case 17, indicating potential swamping. The efficient high breakdown methods Fast-LTS and Fast-S do identify the three outliers as the most outlying cases in the sample, but the residuals are too small to be considered significantly outliers. In contrast, robust methods proposed by Hadi and Simonoff [11] and Fast-PTS estimator identify correctly the clean set $4-25$, with each of the cases $1 - 3$ having residuals greater than 3.78.

data set	n	p	outliers	%outliers identified	time (CPU seconds)
Telephone	24	2	15,16,17,18,19,20	100	0.05
Stars	47	2	11,20,30,34	100	0.18
Wood	20	6	4,6,8,19	100	0.06
Hawkins	75	4	1,2,3,4,5,6,7,8,9,10	100	0.50
Hadi	25	3	1,2,3	100	0.08

Table 2: Performance of the Fast-PTS algorithm on some small data sets.

5.3 Monte Carlo Simulation

To explore further the properties of the PTS method, we have performed an extensive set of simulation experiments for larger sample sizes and observation

dimensions taken from Barrera and Yohai [25]. The experiments compare the performance of our Fast-PTS algorithm with the results of Fast-LTS and Fast-S algorithms.

Barrera and Yohai [25] considered a model as in (1) with $\mathbf{x} = (1, x_1, x_2, \ldots, x_p)$, where a proportion $(1 - \epsilon)$ in the data samples $\mathbf{x} = (x_1, x_2, \ldots, x_p, y)$ follow a multivariate normal distribution. Without loss of generality the mean vector is taken equal to 0 and the identity as the covariance matrix. These observations follow the model in (1) with $\boldsymbol{\beta} = 0$. The contaminated observations are high leverage outliers with $\mathbf{x} = (1, 100, 0, \ldots, 0)$ and $y = \text{slope} \times 100$, where the slope of contamination takes values from 0.9 to 2.0.

Under high leverage contaminations, the objective functions of the three high breakdown estimators typically have two distinct types of local minimum, that is one close to the true value of the regression parameter and a second one close to the slope of the outliers.

	slope											
method	0.9	1.0	1.1	1.2	1.3	1.4	1.5	1.6	1.7	1.8	1.9	2.0
Fast-PTS	58	44	21	8	3	2	1	1	1	0	0	0
Fast-S	88	60	29	8	1	0	0	0	0	0	0	0
Fast-LTS	100	100	95	86	67	42	21	9	4	1	0	0

Table 3: Percentage of samples where convergence occured to the wrong local minimum. $n = 400$, $p = 36$ and 10% of outliers located at $(\mathbf{x}_1, \ldots, \mathbf{x}_{36}, y) = (1, 100, 0, \ldots, 0, \textit{slope} \times 100)$.

Two measures of performance are used in this Monte Carlo study to compare the three estimators:

- The percentage of samples for which each algorithm converged to a wrong local minimum i.e. a local minimum close to the slope of the outliers.
- The mean square error (MSE) of the parameter estimate $\hat{\boldsymbol{\beta}}$.

Several different values of the slope of contamination were used to determine the behaviour of the three algorithms. Its range varies from 0.9 to 2.0 with sample sizes of $n = 400$ and $p = 36$. The proportion of outliers ϵ was set to 10%.

Tables 3 and 4 contain the percentage of convergence to a wrong local minimum and the MSEs respectively. From Table 3 we observe that the Fast-PTS has the lowest wrong convergence with slope between 0.9 and 1.1. For larger slope the results are comparable for all estimators. In Table 4 it is shown that the Fast-PTS has the lowest MSE for the contaminated data with slope between 0.9 and 1.9. For the rest of the data, the MSE of the Fast-PTS is comparable with that of Fast-LTS and Fast-S. The main conclusion from Tables 3 and 4 is that for contamination with small slope, the Fast-PTS perform better than Fast-S and Fast-LTS in percentage of wrong convergence. However, for larger slopes, the PTS has important improvement in the MSEs.

To explore further the comparison of the three estimators, we use again the same data as Barrera and Yohai [25], with different values of n, p, considering

	slope											
method	0.9	1.0	1.1	1.2	1.3	1.4	1.5	1.6	1.7	1.8	1.9	2.0
Fast-PTS	1.54	1.30	0.82	0.48	0.32	0.29	0.36	0.27	0.31	0.30	0.33	0.48
Fast-S	1.65	1.40	0.98	0.62	0.48	0.47	0.45	0.45	0.47	0.46	0.46	0.46
Fast-LTS	1.73	1.46	1.39	1.34	1.08	0.92	0.74	0.58	0.51	0.47	0.46	0.46

Table 4: Mean Squared Errors. $n = 400$, $p = 36$ and 10% of outliers located at $(\mathbf{x}_1, \ldots, \mathbf{x}_{36}, y) = (1, 100, 0, \ldots, 0, slope \times 100)$.

only one value of the contamination slope, $y = 100 \times 1$ and $\epsilon = 10\%$. Tables 5 and 6 show the average time needed until convergence of the algorithm, the percentage of samples where the algorithm converged to a wrong local minimum ("%wrong") and the mean squared error ("MSE") for the three estimators. From Table 5 we observe that the Fast-PTS performs significantly better in both performance criteria especially when the data set is of size $n = 100$ or $n = 500$.

		Fast-PTS			Fast-S		Fast-LTS	
n	p	CPU time	%wrong	MSE	%wrong	MSE	%wrong	MSE
100	2	0.08	0.2	0.030	30.6	0.370	50.8	0.692
	3	0.09	0.2	0.052	37.2	0.506	56.0	0.890
	5	0.12	0.2	0.099	46.8	0.773	64.8	1.193
500	5	2.28	0.0	0.014	15.4	0.193	61.4	0.803
	10	3.78	0.0	0.029	19.8	0.292	65.8	0.935
	20	8.31	0.2	0.069	30.4	0.547	79.4	1.207
1000	5	8.75	0.0	0.007	7.8	0.095	61.2	0.715
	10	14.41	0.0	0.014	5.4	0.089	69.4	0.859
	20	31.11	0.0	0.028	12.4	0.208	82.4	1.158

Table 5: Samples with 10% of outliers located at $(\mathbf{x}_1, \ldots, \mathbf{x}_p, y) = (1, 100, 0, \ldots, 0, 100)$.

For heavier contamination $\epsilon = 0.20$ and slope 2.2, Table 6 shows the results of these simulations. As in Table 5 we observe that the Fast-PTS has the best overall performance especially for the small data sets $n = 100$. However, for larger data sets ($n \geq 500$, the Fast-PTS and Fast-S perform comparable, since the outliers are easier to detect by both robust estimators due to the large slope value.

To continue the implementation of the Fast-PTS algorithm and compare it with the well-known methods discussed in this article, we perform extra Monte Carlo experiments to evaluate the performance of our robust procedure. To carry out one simulation run, we proceeded as follows. The distributions of independent variables and errors and the values of parameters are given. The observations y_i, were obtained following the regression model as in (1) with $p = 3$ and $n = 50$, where the coefficient values are $\beta_1 = 1.20$, $\beta_2 = -0.80$ and a zero constant term $\beta_0 = 0.0$. We prefer the Gauss distribution for the iid error term $u \sim N(0, \sigma^2 = 16^2)$, while $\mathbf{x}_1$ and x_2 are iid values drawn also from normal distributions $N(\mu = 20, \sigma^2 = 6^2)$ and $N(\mu = 30, \sigma^2 = 8^2)$ respectively.

		Fast-PTS			Fast-S		Fast-LTS	
n	p	CPU time	%wrong	MSE	%wrong	MSE	%wrong	MSE
100	2	0.08	0.2	0.032	10.0	0.543	40.9	2.333
	3	0.09	0.0	0.056	15.6	0.947	51.6	3.218
	5	0.11	0.8	0.109	34.0	2.303	70.2	5.241
500	5	2.28	0.0	0.016	0.0	0.025	23.4	1.236
	10	3.74	0.0	0.032	0.0	0.076	33.8	2.102
	20	7.82	11.2	0.873	9.2	0.734	72.2	4.276
1000	5	9.40	0.0	0.008	0.0	0.012	11.0	0.542
	10	14.94	0.0	0.016	0.0	0.025	20.4	1.205
	20	31.63	0.0	0.030	0.0	0.055	49.8	2.980

Table 6: Samples with 20% of outliers located at $(\mathbf{x}_1, \ldots, \mathbf{x}_p, y) = (1, 100, 0, \ldots, 0, 220)$.

We consider that the sample may contain three types of outliers, regression outliers ("bad" high-leverage points), "good" high-leverage points, and response outliers (y-outliers). An extra value is drawn from the uniform distribution $U(a = 80, b = 220)$ and for the regression outlier is added to x_1 or x_2, for the "good" leverage point is added to x_1 or x_2 but the value of the dependent variable y follows their contamination, according to the above regression model, for the response outlier is added to y. All simulation results are based on 150 replications enough to obtain a relative error $< 10\%$ with a reasonable confidence level of at least 90% for all the simulation estimates. The robust scale estimate $\hat{\sigma}$ from LTS with coverage $k = 28$ is used throughout the simulation study.

Tables 7 and 8 display the results concerning the performance of the three robust estimators corresponding to the following cases: Table 7 is based on data contaminated by "bad" and "good" high leverage points whereas Table 8 is based on data contaminated by "bad" high leverage outliers (heavier contamination).

	Fast-PTS	Fast-S	Fast-LTS
%wrong	2.7	13.3	16.0
mean estimate of $\hat{\beta}_0$	-1.796	-1.633	-2.782
mean estimate of $\hat{\beta}_1$	1.176	1.129	1.147
mean estimate of $\hat{\beta}_2$	-0.766	-0.758	-0.732
mean squared error of $\hat{\beta}_0$	35.38	58.41	112.21
mean squared error of $\hat{\beta}_1$	0.036	0.106	0.132
mean squared error of $\hat{\beta}_2$	0.015	0.030	0.049
computation time (secs)	0.56	0.29	0.18

Table 7: Samples with 32% of outliers (x-outliers=6, "good" leverage points=4, y-outliers=6), $n = 50$, $p = 3$. True: $\beta_0 = 0.0$, $\beta_1 = 1.20$, $\beta_2 = -0.80$.

6 Conclusions and Future Research

The PTS estimator can be considered as a procedure which instead of improving upon previous robust estimators such as LTS and MCD, it combines the infor-

	Fast-PTS	Fast-S	Fast-LTS
%wrong	26.0	61.3	50.0
mean estimate of $\hat{\beta}_0$	-3.955	-1.607	-2.123
mean estimate of $\hat{\beta}_1$	1.131	0.814	0.910
mean estimate of $\hat{\beta}_2$	-0.646	-0.569	-0.600
mean squared error of $\hat{\beta}_0$	84.41	158.25	173.98
mean squared error of $\hat{\beta}_1$	0.088	0.473	0.390
mean squared error of $\hat{\beta}_2$	0.084	0.167	0.154
CPU time	0.55	0.27	0.16

Table 8: Samples with 32% of outliers (x-outliers=10, y-outliers=6), $n = 50$, $p = 3$. True: $\beta_0 = 0.0$, $\beta_1 = 1.20$, $\beta_2 = -0.80$.

mation so obtained from them in a penalized manner, so as to simultaneously retain their favorable robust characteristics, while being efficient and identifying multiple masked outliers. Specifically, this is achieved by using a robust scale and leverage from the LTS and MCD respectively. Moreover, an efficient fast algorithm which is based on a set of necessary conditions for local optimality is also presented. Extensive computational experiments on a large set of instances, indicate that the proposed estimator ouperforms other robust estimators in the presence of all possible types of outliers.

Future research can be performed on both the estimator and the algorithm. With respect to the estimator, a different robust scale $\hat{\sigma}$ could be investigated, and also it can be extended to treat data sets which contain a mixture of multilinear regression models. The algorithm could also be modified to handle massive data sets, which may have millions of observations, by utilizing the natural parallel structure of the algorithm.

References

[1] C. Agostinelli and M. Markatou. A one-step robust estimator for regression based on the weighted likelihood reweighting scheme. *Statistics and Probability Letters*, 37:342–350, 1998.

[2] J. Agulló. New algorithms for computing the least trimmed squares regression estimator. *Computational Statistics and Data Analysis*, 36:425–439, 2001.

[3] A.C. Atkinson. Fast very robust methods for the detection of multiple outliers. *Journal of the American Statistical Association*, 89:1329–1339, 1994.

[4] A.C. Atkinson and T.-C. Cheng. Computing least trimmed squares regression with the forward search. *Statistics and Computing*, 9:251–263, 1999.

[5] M.S. Bazaraa, H.D. Sherali, and C.M. Shetty. *Nonlinear Programming: Theory and Algorithms*. Wiley, 1993.

[6] C.W. Coakley and T.P. Hettmansperger. A bounded influence, high breakdown, efficient regression estimator. *Journal of the American Statistical Association*, 88:872–880, 1993.

[7] D.L. Donoho and P.J. Huber. The notion of breakdown point. In *A Festschrift for Erich Lehmann*, Belmont, CA, 1983. Wadsworth.

[8] J.F. Gentleman and M.B. Wilk. Detecting outliers ii: Supplementing the direct analysis of residuals. *Biometrics*, 31:387–410, 1975.

[9] D. Gervini and V.J. Yohai. A class of robust and fully efficient regression estimators. *Annals of Statistics*, 30:583–616, 2002.

[10] A. Giloni and M. Padberg. Least trimmed squares regression, least median squares regression, and mathematical programming. *Mathematical and Computer Modelling*, 35:1043–1060, 2002.

[11] A. S. Hadi and J.S. Simonoff. Procedures for the identification of multiple outliers in linear models. *Journal of the American Statistical Association*, 88:1264–1272, 1993.

[12] D.M. Hawkins. The feasible solution algorithm for least trimmed squares regression. *Computational Statistics and Data Analysis*, 17:185–196, 1994.

[13] D.M. Hawkins, D. Bradu, and G.V. Kass. Location of several outliers in multiple regression data using elemental sets. *Technometrics*, 26:197–208, 1984.

[14] D.M. Hawkins and D.J. Olive. Improved feasible solution algorithms for high breakdown estimation. *Computational Statistics and Data Analysis*, 30:1–11, 1999.

[15] O. Hössjer. Exact computation of the least trimmed squares estimate in simple linear regression. *Computational Statistics and Data Analysis*, 19:265–282, 1995.

[16] L.M. Li. An algorithm for computing exact least-trimmed squares estimate of simple linear regression with constraints. *Computational Statistics and Data Analysis*, 48:717–734, 2005.

[17] D. Peña and V.J. Yohai. A fast procedure for outlier diagnostics in large regression problems. *Journal of the American Statistical Association*, 94:434–445, 1999.

[18] L.S. Pitsoulis and M.G.C. Resende. Greedy randomized adaptive search procedures. In P.M. Pardalos and M.G.C. Resende, editors, *Handbook of Applied Optimization*, pages 168–183. Oxford University Press, 2002.

[19] P.J. Rousseeuw. Least median of squares regression. *Journal of the American Statistical Association*, 79:871–880, 1984.

[20] P.J. Rousseeuw and K. Van Driessen. A fast algorithm for the minimum covariance determinant estimator. *Technometrics*, 41:212–223, 1999.

[21] P.J. Rousseeuw and K. Van Driessen. Computing LTS regression for large data sets. *Data Mining and Knowledge Discovery*, 12:29–45, 2006.

[22] P.J. Rousseeuw and A.M. Leroy. *Robust Regression and Outlier Detection*. John Wiley, New York, 1987.

[23] P.J. Rousseeuw and V.J. Yohai. Robust regression by means of s-estimators. In *Robust and Nonlinear Time Series Analysis*, pages 256–272. Springer Verlag, 1984.

[24] R.J. Rousseeuw and B.C. Van Zomeren. Unmasking multivariate outliers and leverage points. *Journal of the American Statistical Association*, 85:633–639, 1990.

[25] M. Salibian-Barrera and V.J. Yohai. A fast algorithm for s-regression estimates. *Journal of Computational and Graphical Statistics*, 15(2):414–427, 2006.

[26] V.J. Yohai. High breakdown-point and high efficiency robust estimates for regression. *Annals of Statistics*, 15:642–656, 1987.

[27] V.J. Yohai and R.H. Zamar. High breakdown point estimates of regression by means of the minimization of an efficient scale. *Journal of the American Statistical Association*, 83:406–413, 1988.

[28] G. Zioutas and A. Avramidis. Deleting outliers in robust regression with mixed integer programming. *Acta Mathematicae Applicatae Sinica*, 21:323–334, 2005.

[29] G. Zioutas, A. Avramidis, and L. Pitsoulis. Penalized trimmed squares and a modification of support vectors for unmasking outliers in linear regression. *REVSTAT*, 5:115–136, 2007.